Learning Domain-Independent Policies for Open List Selection

André Biedenkapp, David Speck, Silvan Sievers,
Frank Hutter, Marius Lindauer, Jendrik Seipp

Motivation

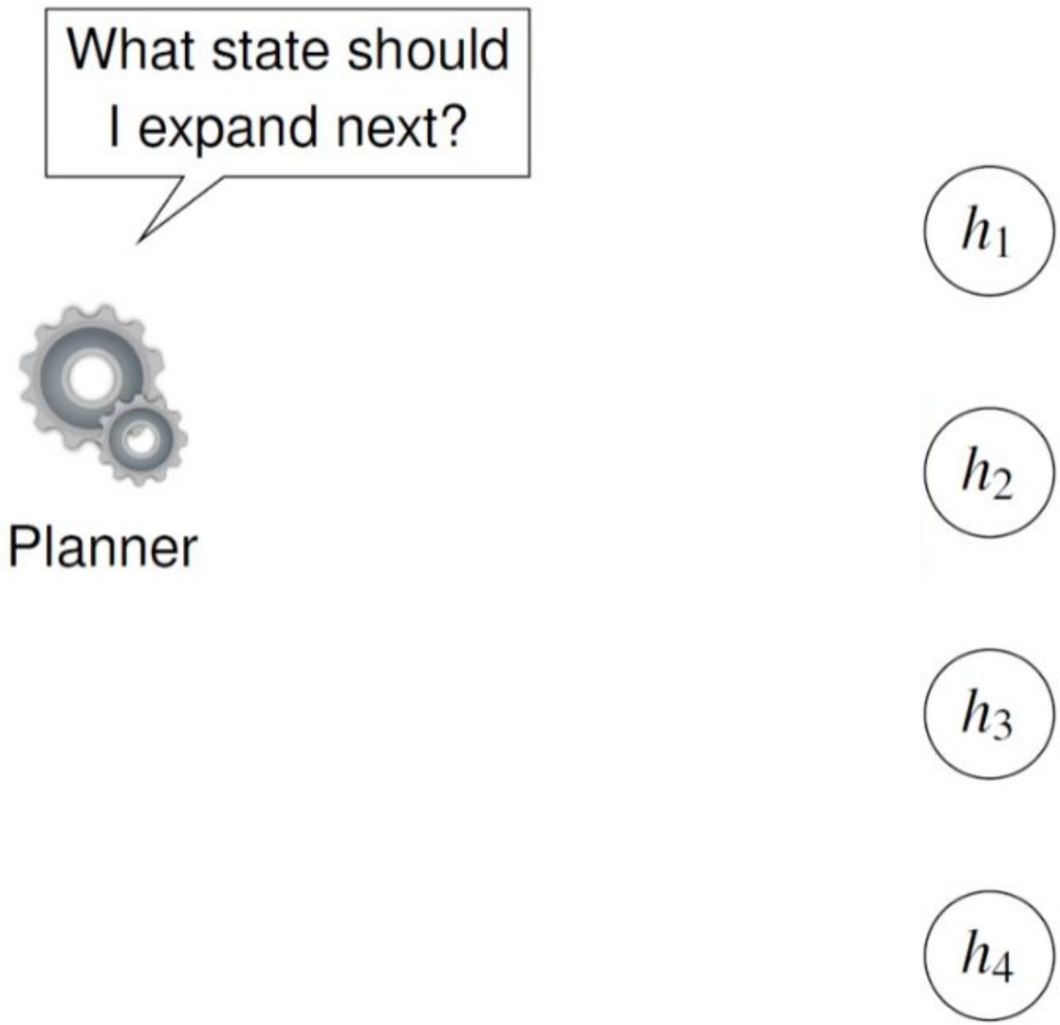

Motivation

Satisficing planning with multiple heuristics

- Search for a good plan
- Inadmissible heuristics are difficult to combine
- States evaluated with each heuristic
- One separate open list for each heuristic

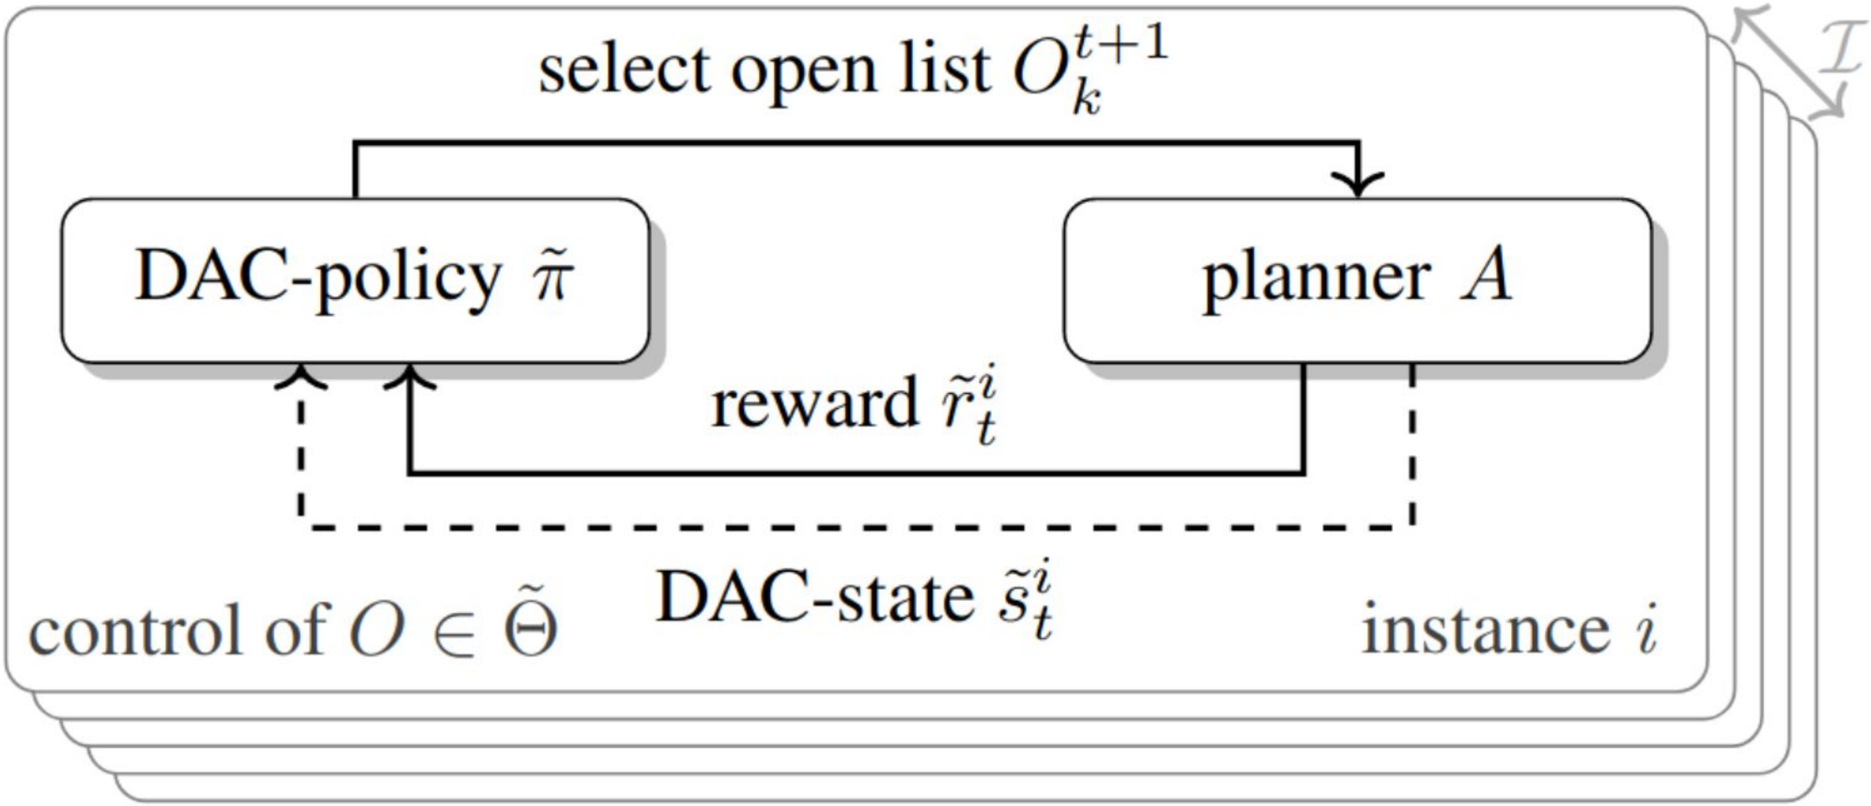

Speck, D.; Biedenkapp, A.; Hutter, F.; Mattmüller, R.; and Lindauer, M. 2021. Learning Heuristic Selection with Dynamic Algorithm Configuration. In Proc. ICAPS 2021, 597–605.

Is Domain Dependent Good enough?

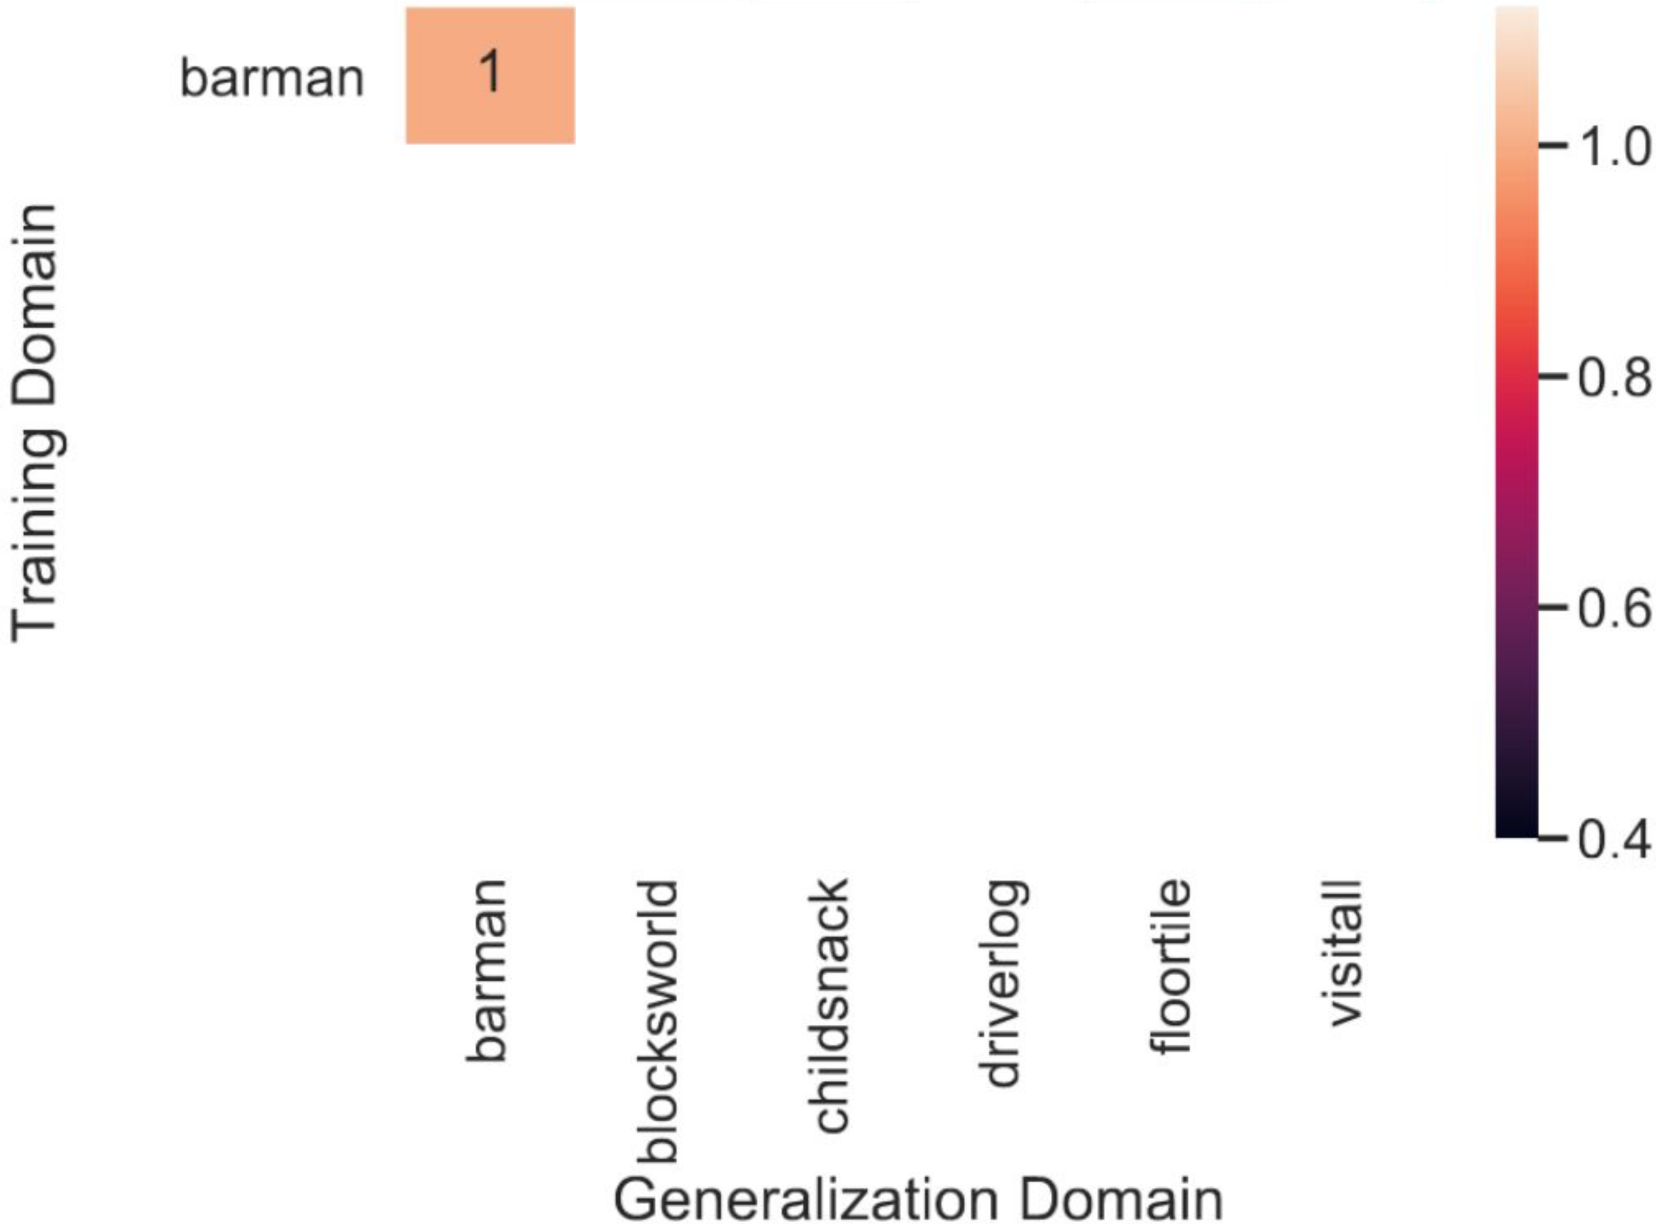

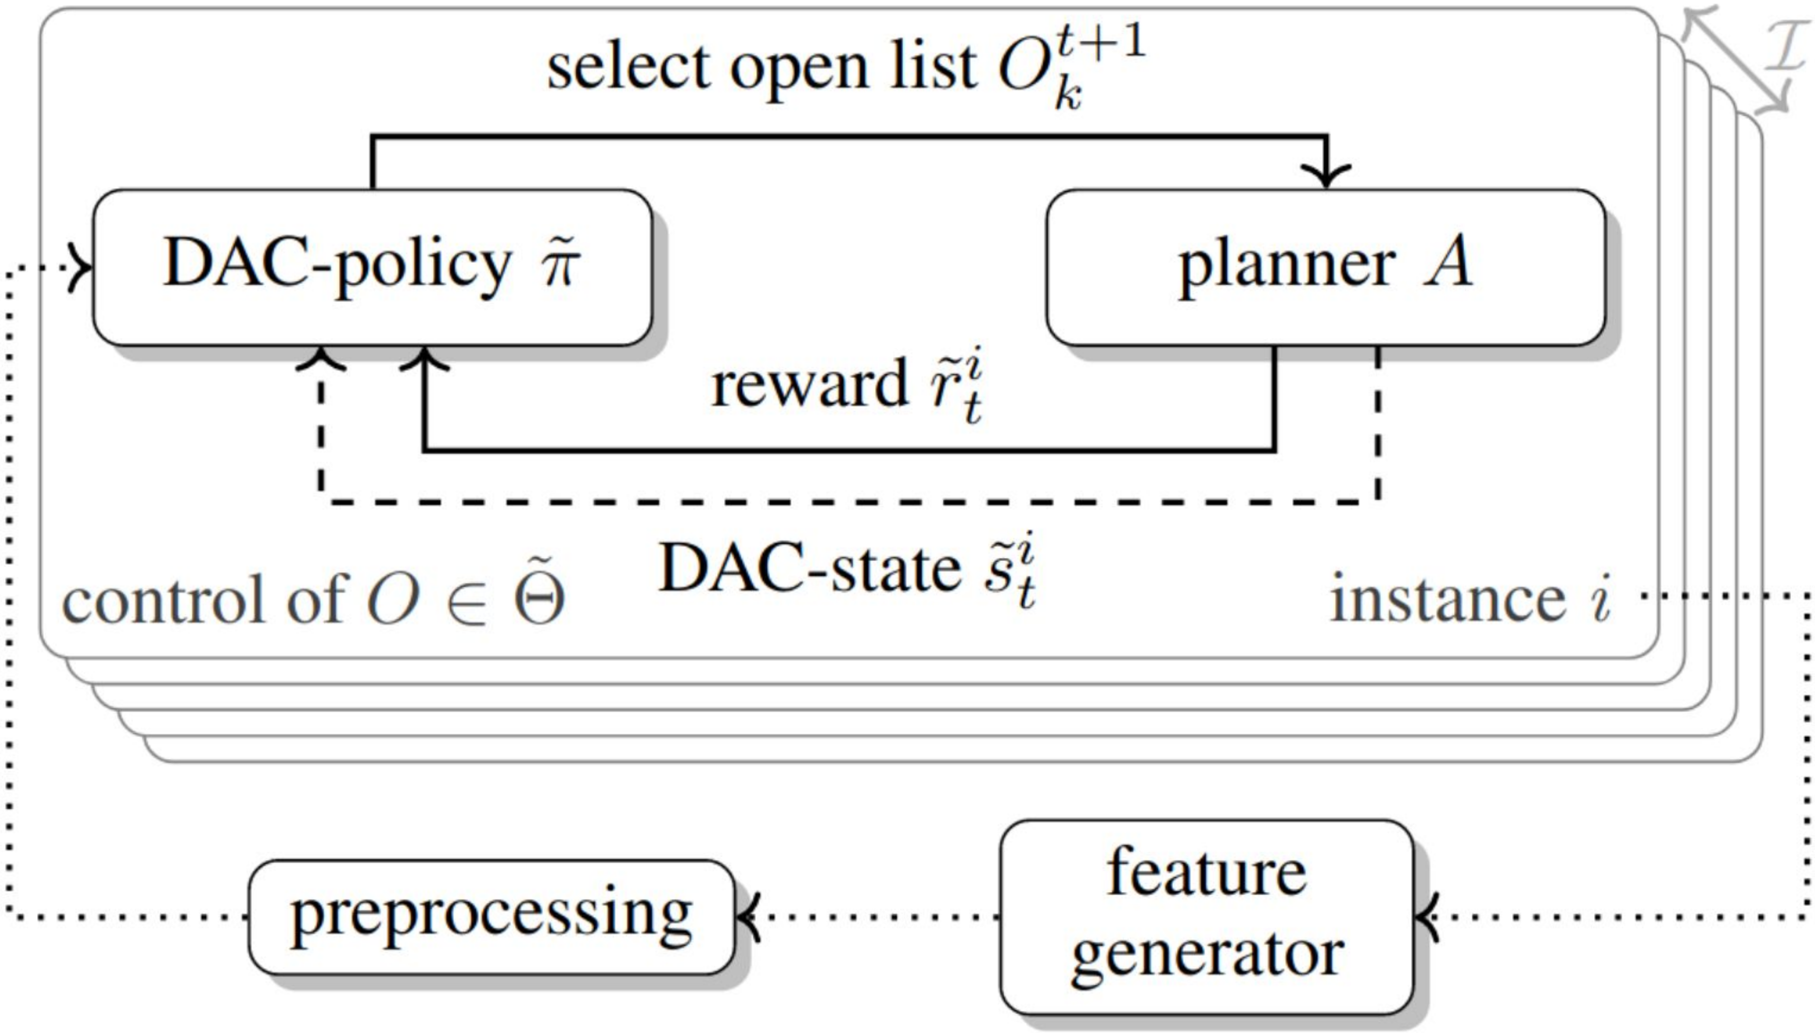
select open list O_k^{t+1}
DAC-policy $\tilde{\pi}$
planner A
reward $\tilde{r}_t^i$
DAC-state $\tilde{s}_t^i$
control of $O \in \tilde{\Theta}$
instance i
$\mathcal{I}$
preprocessing
feature generator

- Using Only Dynamic Features (no-F)
- Concatenating Instance and Dynamic Features (raw-F)
- Learning Separate Representations (embed)
- Decoupling Instance and Dynamic Features (dc)

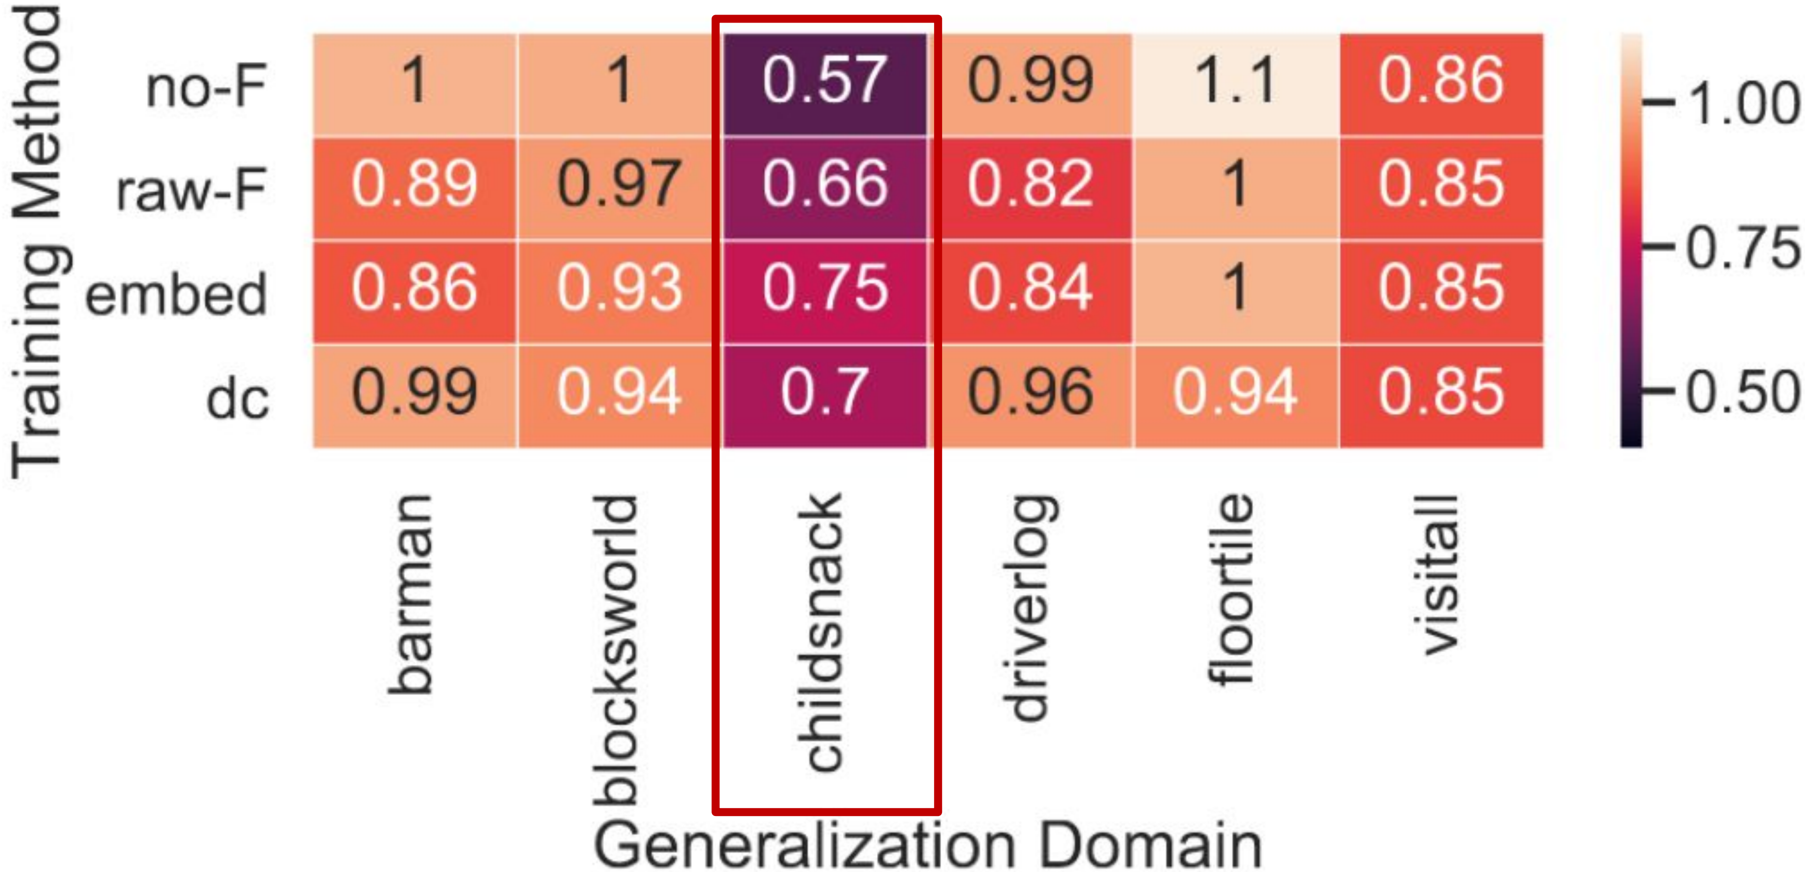
Training Method
no-F
raw-F
embed
dc
1
1
0.57
0.99
1.1
0.86
0.89
0.97
0.66
0.82
1
0.85
0.86
0.93
0.75
0.84
1
0.85
0.99
0.94
0.7
0.96
0.94
0.85
1.00
0.75
0.50
barman
blocksworld
childsnack
driverlog
floortile
visitall
Generalization Domain

Looking forward to meeting you at the poster!

Motivation - DAC

Goal: Select configuration for the problem at hand and adapt while planning

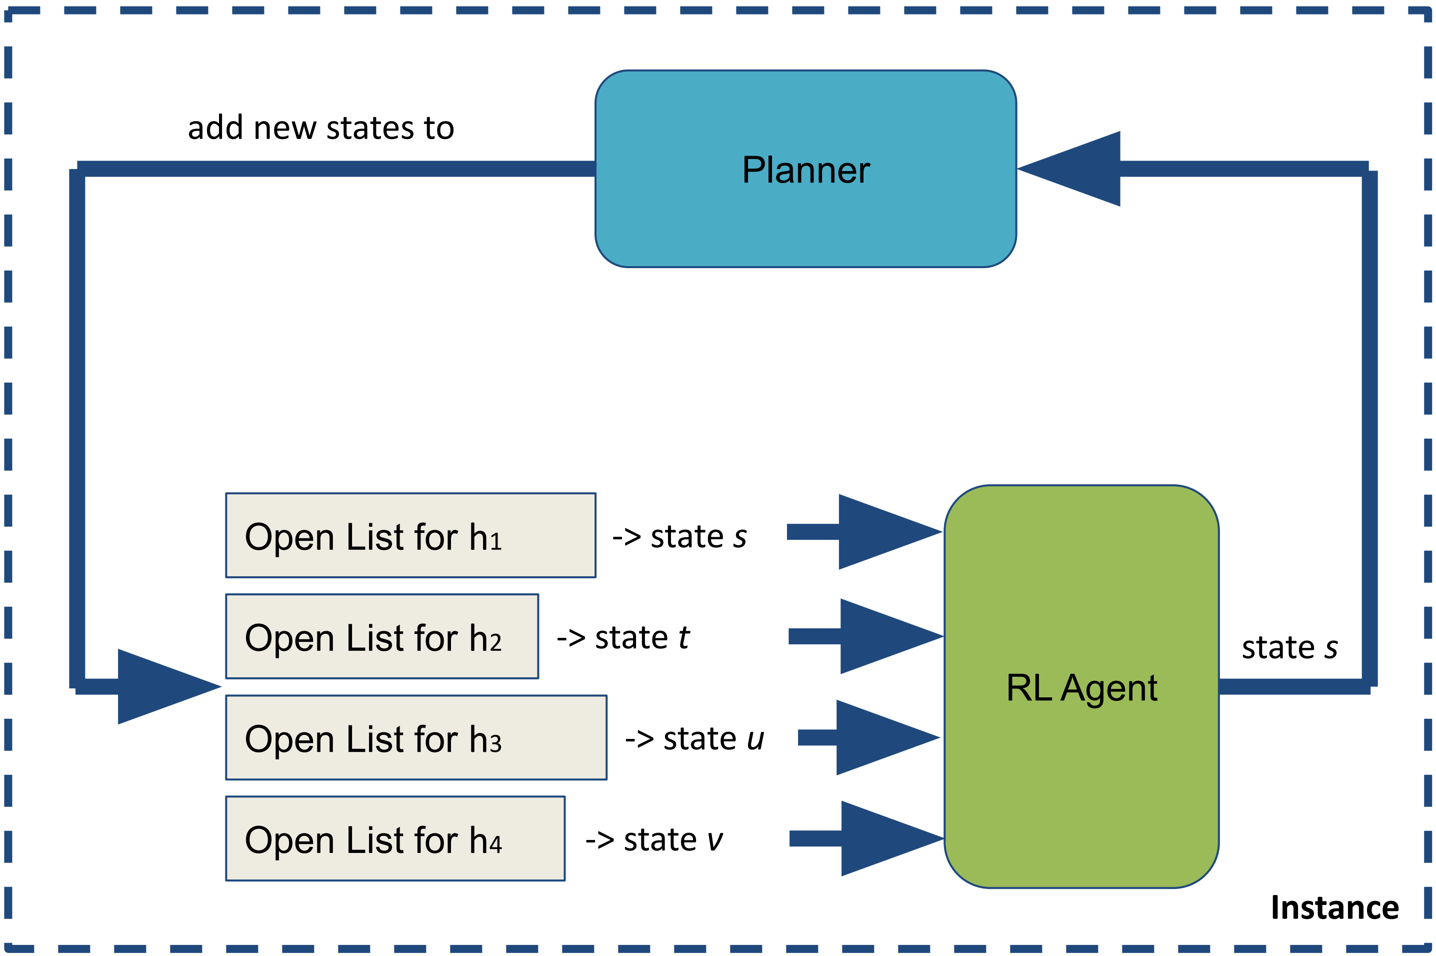

Motivation

Configuration -> Planner -> Performance

Motivation - AC

Goal: Find single best configuration

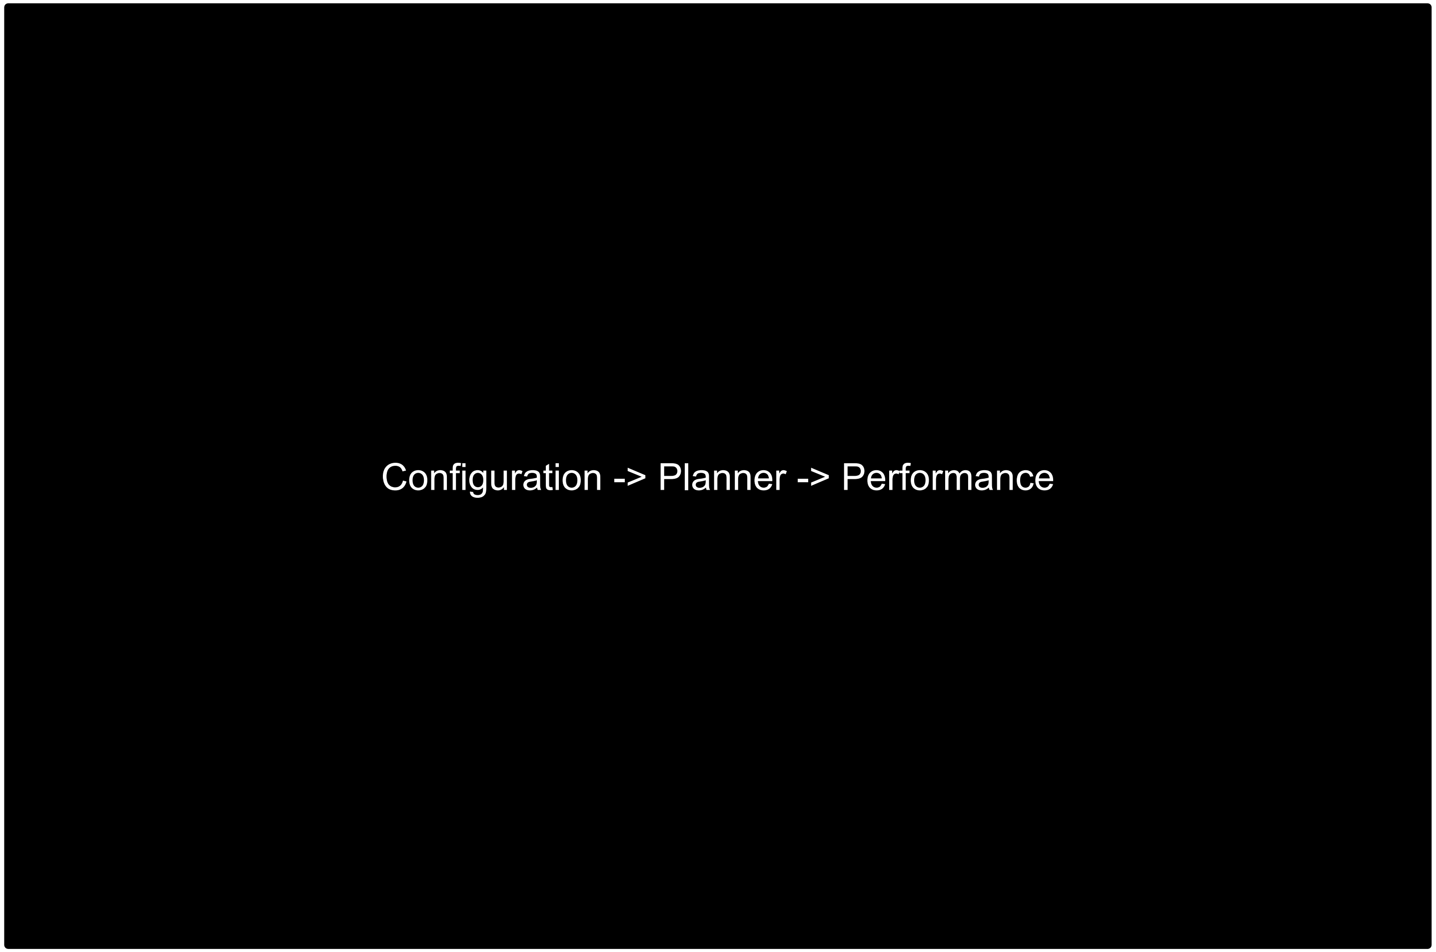

Motivation - AS

Goal: Select best configuration for the problem at hand

Configuration
+
Instance Features

-> Planner -> Performance

Motivation - DAC

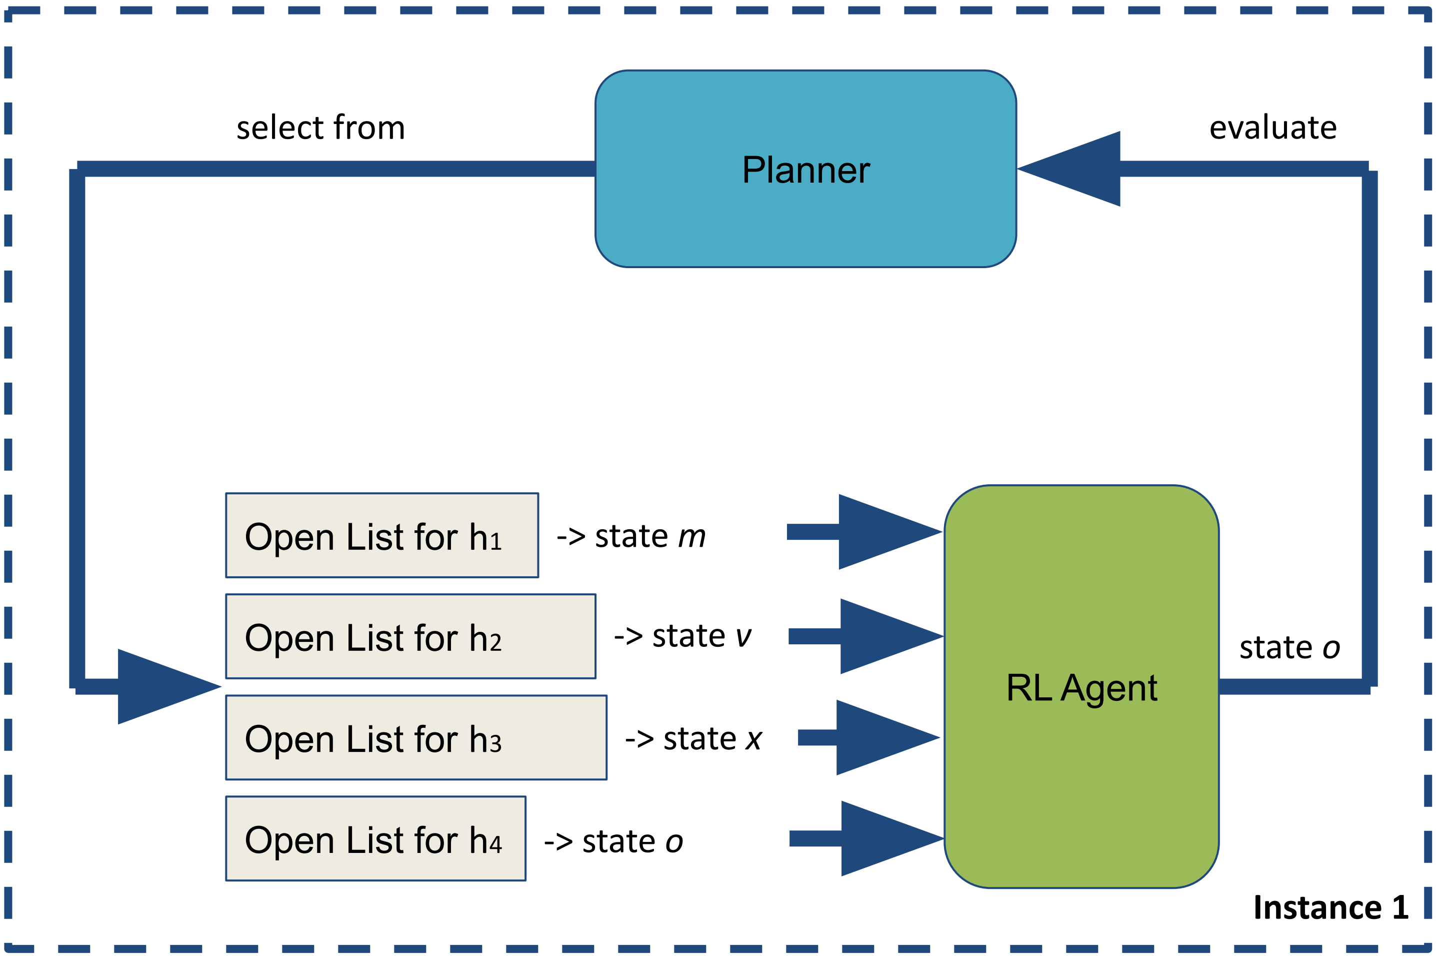

Motivation - DAC

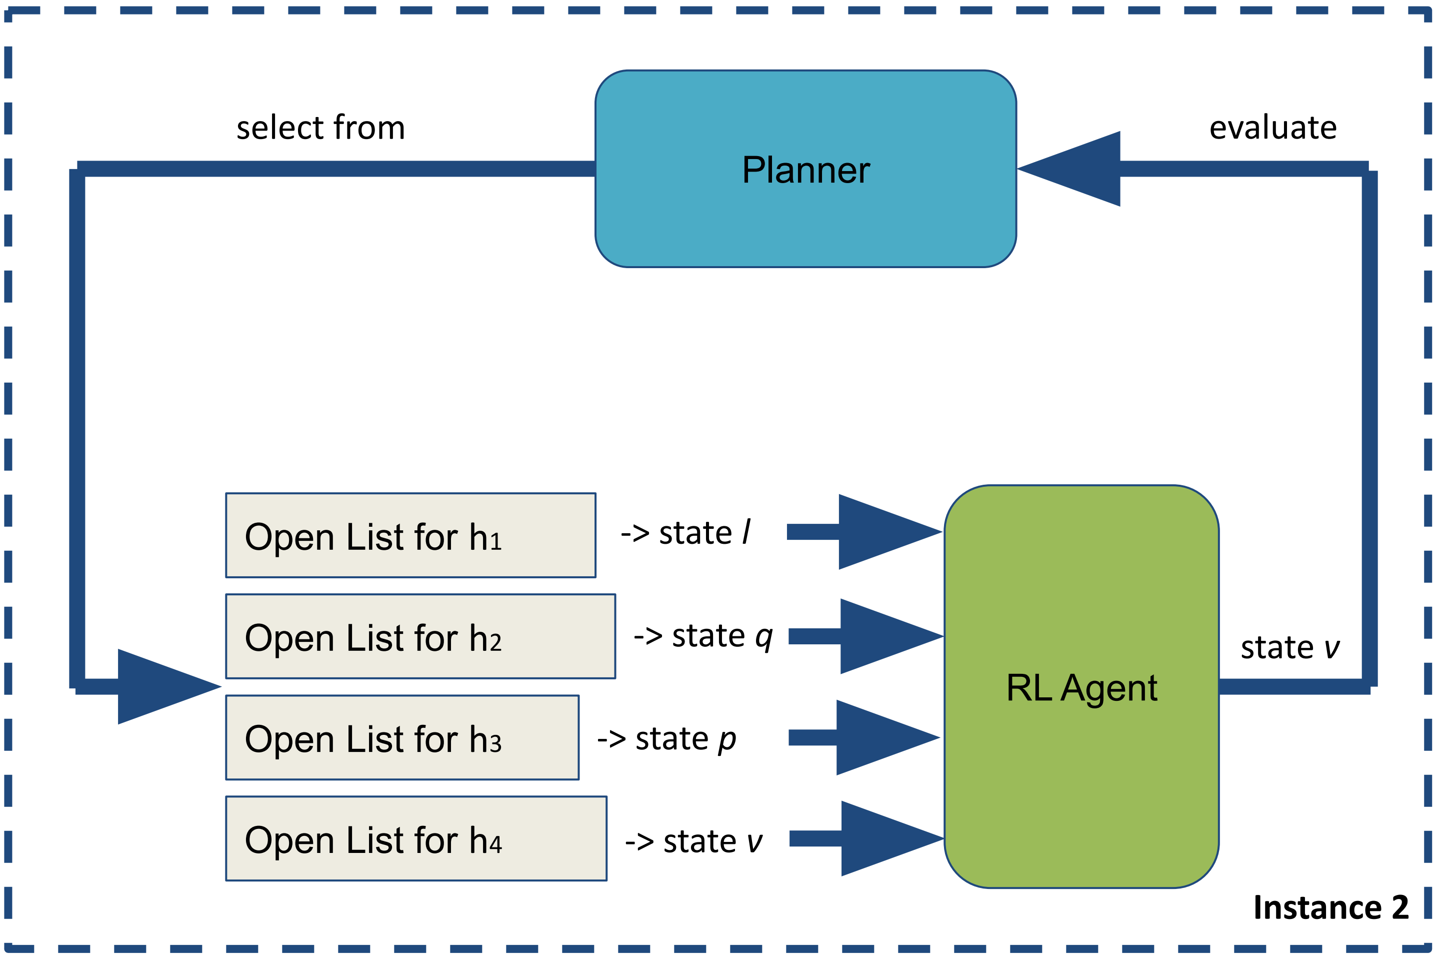

In a Nutshell

1. We generalize previous DAC approaches to learn domain-independent open list selection policies
2. We present novel ways to learn from instance specific features jointly with dynamic features
3. Our learned policies reduce the required number of node expansions on several domains
4. We use DAC as a tool to gain insights on why LAMA's policy has such strong performance